Refine Search

Search Results -

Terms	Documents
L3 and (processor near10 cache)	87

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search: L4

Refine Search | Recall Text | Clear | Interrupt

Search History

DATE: Tuesday, November 16, 2004 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=PGPB,USPT,USOC; PLUR=YES; OP=OR			
L4	L3 and (processor near10 cache)	87	L4
L3	L1 and (add$3 same request same node)	123	L3
L2	L1 and (add$3 same node)	423	L2
L1	((shared or common) near5 memory) same (multiprocessor or (multi adj1 processor))	3138	L1

END OF SEARCH HISTORY

h e b b cg b e e ch

Refine Search

Search Results -

Terms	Documents
L4	0

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search: L5

Refine Search | Recall Text | Clear | Interrupt

Search History

DATE: Tuesday, November 16, 2004 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR			
L5	L4	0	L5
DB=PGPB,USPT,USOC; PLUR=YES; OP=OR			
L4	L3 and (processor near10 cache)	87	L4
L3	L1 and (add$3 same request same node)	123	L3
L2	L1 and (add$3 same node)	423	L2
L1	((shared or common) near5 memory) same (multiprocessor or (multi adj1 processor))	3138	L1

END OF SEARCH HISTORY

h e b b cg b e e ch

Refine Search

Search Results -

Terms	Documents
(700/5 |709/213 |709/214 |709/251 |710/305 |710/317 |710/300 |710/62 |710/4 |710/72 |711/141 |711/148 |711/120 |712/14 |712/211).ccls.	5971

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search: L6

Refine Search | Recall Text | Clear | Interrupt

Search History

DATE: Tuesday, November 16, 2004 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
	DB=PGPB,USPT,USOC; PLUR=YES; OP=OR		
L6	710/305,317,300,62,4,72;711/141,148,120;709/213,214,251;700/5;712/14,211.ccls.	5971	L6
	DB=EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR		
L5	L4	0	L5
	DB=PGPB,USPT,USOC; PLUR=YES; OP=OR		
L4	L3 and (processor near10 cache)	87	L4
L3	L1 and (add$3 same request same node)	123	L3
L2	L1 and (add$3 same node)	423	L2
L1	((shared or common) near5 memory) same (multiprocessor or (multi adj1 processor))	3138	L1

END OF SEARCH HISTORY

h e b b cg b e e ch

Refine Search

Search Results -

Terms	Documents
L4 and L6	49

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search: L7

Refine Search | Recall Text | Clear | Interrupt

Search History

DATE: Tuesday, November 16, 2004 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
	DB=PGPB,USPT,USOC; PLUR=YES; OP=OR		
L7	l4 and L6	49	L7
L6	710/305,317,300,62,4,72;711/141,148,120;709/213,214,251;700/5;712/14,211.ccls.	5971	L6
	DB=EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR		
L5	L4	0	L5
	DB=PGPB,USPT,USOC; PLUR=YES; OP=OR		
L4	L3 and (processor near10 cache)	87	L4
L3	L1 and (add$3 same request same node)	123	L3
L2	L1 and (add$3 same node)	423	L2
L1	((shared or common) near5 memory) same (multiprocessor or (multi adj1 processor))	3138	L1

END OF SEARCH HISTORY

h e b b cg b e e ch

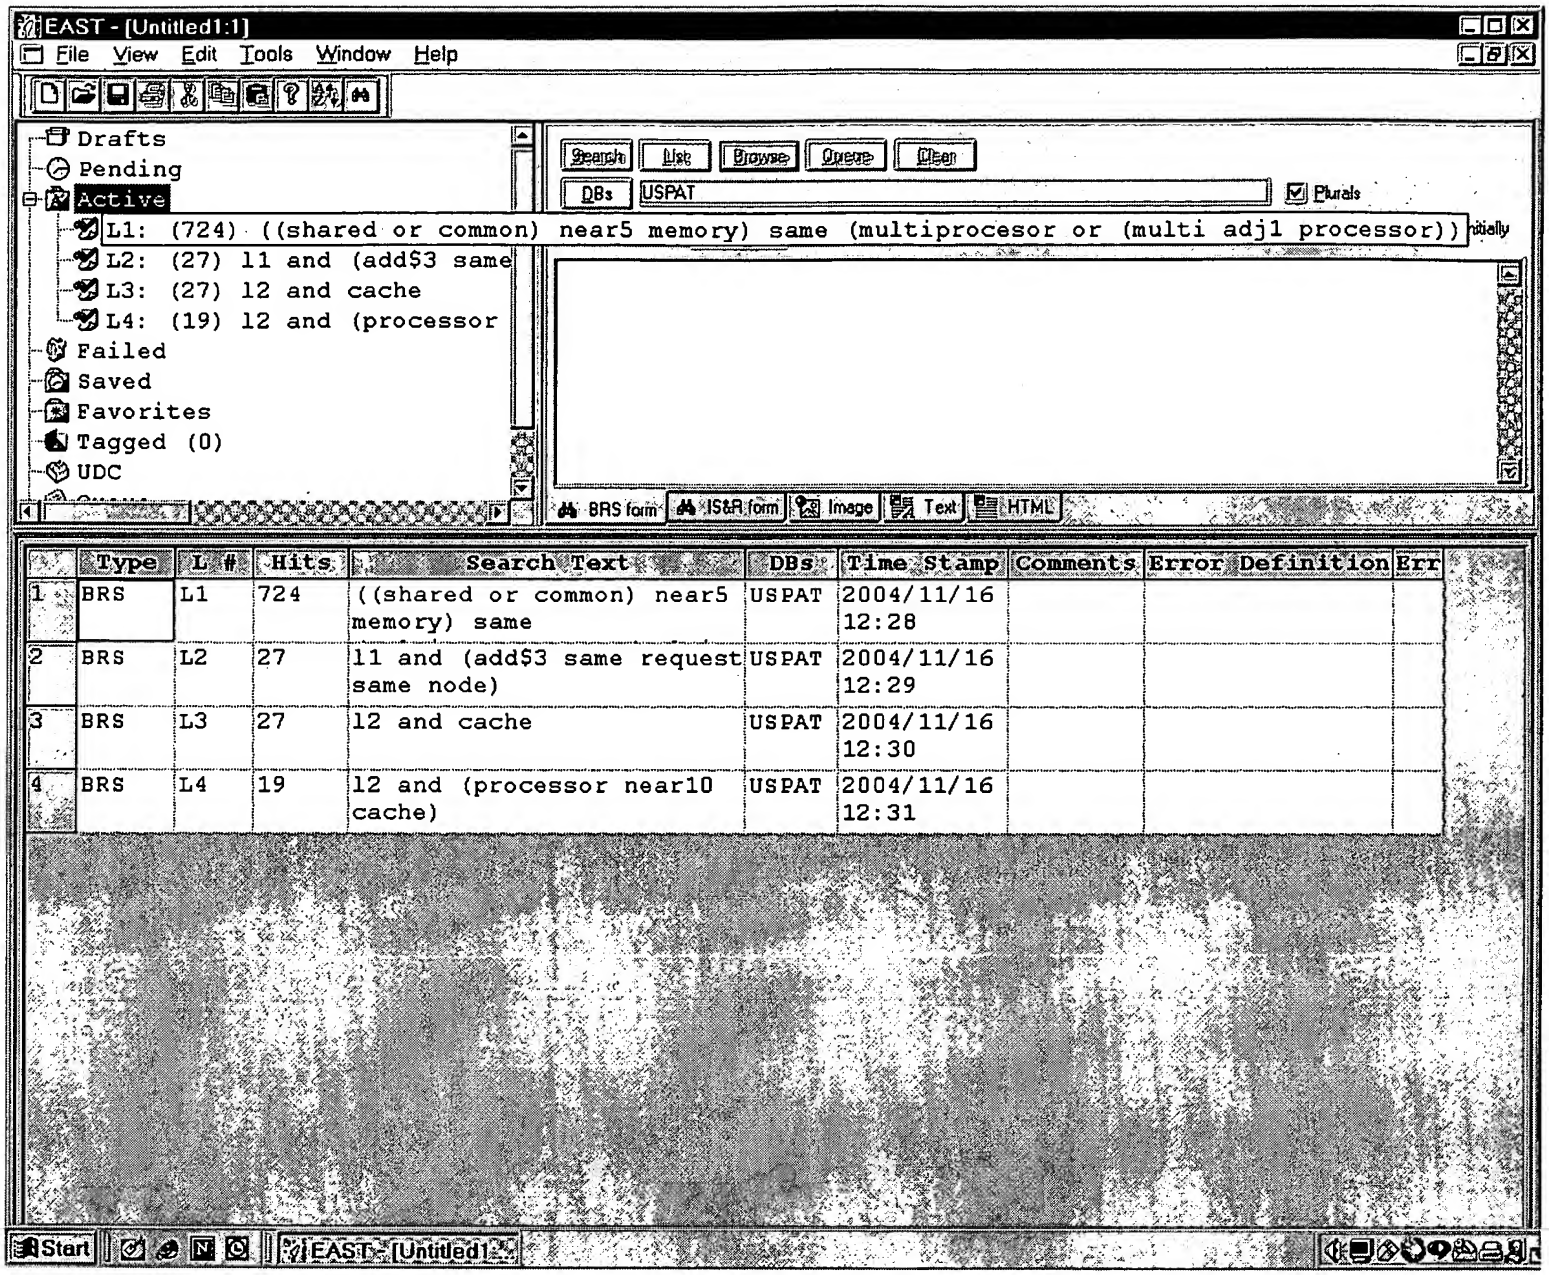

	Type	L #	Hits	Search Text	DBs	Time Stamp	Comments	Error Definition	Err
1	BRS	L1	724	((shared or common) near5 memory) same	USPAT	2004/11/16 12:28			
2	BRS	L2	27	l1 and (add$3 same request same node)	USPAT	2004/11/16 12:29			
3	BRS	L3	27	l2 and cache	USPAT	2004/11/16 12:30			
4	BRS	L4	19	l2 and (processor near10 cache)	USPAT	2004/11/16 12:31			

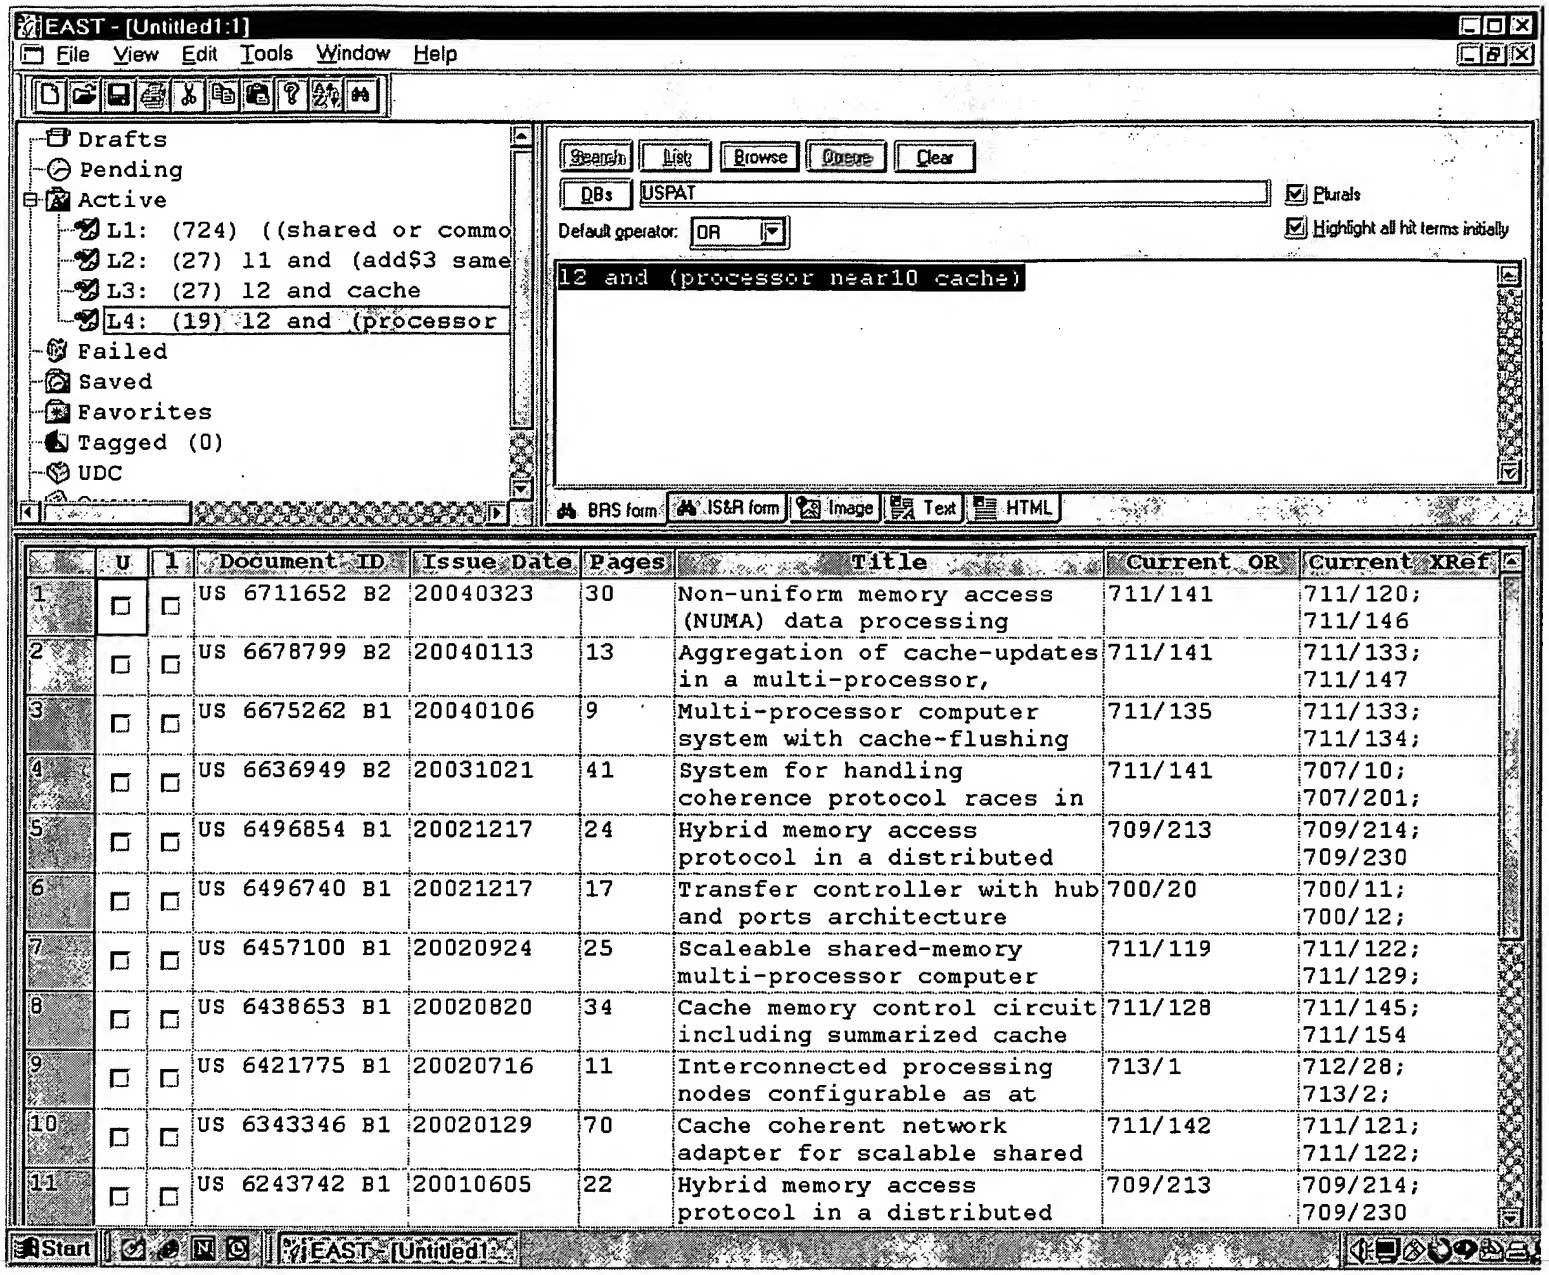

	U	1	Document ID	Issue Date	Pages	Title	Current OR	Current XRef
1	☐	☐	US 6711652 B2	20040323	30	Non-uniform memory access (NUMA) data processing	711/141	711/120; 711/146
2	☐	☐	US 6678799 B2	20040113	13	Aggregation of cache-updates in a multi-processor,	711/141	711/133; 711/147
3	☐	☐	US 6675262 B1	20040106	9	Multi-processor computer system with cache-flushing	711/135	711/133; 711/134;
4	☐	☐	US 6636949 B2	20031021	41	System for handling coherence protocol races in	711/141	707/10; 707/201;
5	☐	☐	US 6496854 B1	20021217	24	Hybrid memory access protocol in a distributed	709/213	709/214; 709/230
6	☐	☐	US 6496740 B1	20021217	17	Transfer controller with hub and ports architecture	700/20	700/11; 700/12;
7	☐	☐	US 6457100 B1	20020924	25	Scaleable shared-memory multi-processor computer	711/119	711/122; 711/129;
8	☐	☐	US 6438653 B1	20020820	34	Cache memory control circuit including summarized cache	711/128	711/145; 711/154
9	☐	☐	US 6421775 B1	20020716	11	Interconnected processing nodes configurable as at	713/1	712/28; 713/2;
10	☐	☐	US 6343346 B1	20020129	70	Cache coherent network adapter for scalable shared	711/142	711/121; 711/122;
11	☐	☐	US 6243742 B1	20010605	22	Hybrid memory access protocol in a distributed	709/213	709/214; 709/230

Your search matched **2** of **1094442** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or enteri new one in the text box.

(shared or common) and memory and processor an Search

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

---

1 **The effect of limited network bandwidth and its utilization by latenc hiding techniques in large-scale shared memory systems**
Sunil Kim; Veidenbaum, A.V.;
Parallel Architectures and Compilation Techniques., 1997. Proceedings. 1997 International Conference on , 10-14 Nov. 1997
Pages:40 - 51

[Abstract] [PDF Full-Text (1284 KB)] **IEEE CNF**

---

2 **Whoops!: a clustered Web cache for DSM systems using memory ma networks**
Cecchet, E.;
Distributed Computing Systems Workshops, 2002. Proceedings. 22nd Internal Conference on , 2-5 July 2002
Pages:806 - 811

[Abstract] [PDF Full-Text (431 KB)] **IEEE CNF**

---

h eee e eee g e ch e ch e e c e e e c e

Whoops!: a clustered Web cache for DSM systems u memory mapped networks

Cecchet, E.

This paper appears in: **Distributed Computing Systems Workshops, 200: Proceedings. 22nd International Conference on**

Publication Date: 2-5 July 2002
On page(s): 806 - 811
ISSN:
Number of Pages: xxviii+835
Inspec Accession Number: 7432563

Abstract:
We present Whoops!, a clustered Web **cache** prototype based on SciFS, a dis **shared memory** (DSM) that benefits from the high performances and the re **addressing** capabilities of **memory** mapped networks like Scalable Coherent (SCI). Whoops! uses the DSM for all Web **cache** management and **cache** sto a **memory** mapped network and a DSM programming model allows us to inve new algorithm to distribute and handle **requests**. We present a new impleme TCP handoff that directly maps remote TCP/IP stacks through the network. Th technique reduces **processor** overhead and forwards TCP acknowledgements microseconds. We have also designed parallel pull-based LRU (PPBL), an effic **request** distribution algorithm for use with DSM systems. The decision is dist all **nodes** thus providing better scalability. PPBL supports multi-frontend envi letting the DSM handle data distribution. Finally, Whoops! implements on the compression when fetching documents from the Web and on the fly decompre sending documents to clients. We show how this technique can reduce paging the DSM and improve overall **cache** performance.

Index Terms:
Internet **cache** storage client-server systems data compression distributed **shared** systems document handling network servers system buses transport protocols DS Scalable Coherent Interface SciFS TCP acknowledgements TCP handoff TCP/IP s **cache** management Whoops! **cache** storage clustered Web **cache** compression distribution decompression distributed **shared memory** **memory** mapped networks

h eee e eee g e ch ch b c be be

frontend environments parallel pull-based LRU **processor** overhead remote **address** capabilities **request** distribution algorithm scalability

Documents that cite this document
There are no citing documents available in IEEE Xplore at this time.

Search Results [PDF FULL-TEXT 431 KB] PREV DOWNLOAD CITATION

h eee e eee g e ch ch b c be be

L7: Entry 1 of 49 File: PGPB Oct 28, 2004

PGPUB-DOCUMENT-NUMBER: 20040215895
PGPUB-FILING-TYPE: new
DOCUMENT-IDENTIFIER: US 20040215895 A1

TITLE: Multi-node computer system in which networks in different nodes implement different conveyance modes

PUBLICATION-DATE: October 28, 2004

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-47
Cypher, Robert E.	Saratoga	CA	US	

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	COUNTRY	TYPE CODE
Sun Microsystems, Inc.				02

APPL-NO: 10/ 813891 [PALM]
DATE FILED: March 31, 2004

RELATED-US-APPL-DATA:
Application is a non-provisional-of-provisional application 60/461997, filed April 11, 2003,

INT-CL: [07] G06 F 12/00, H04 L 12/28, G06 F 13/00

US-CL-PUBLISHED: 711/141; 370/390
US-CL-CURRENT: 711/141; 370/390

REPRESENTATIVE-FIGURES: 20

ABSTRACT:

A system may include several nodes coupled by in inter-node network. Each node includes several active devices coupled by an address network. The address network included in one of the nodes may be configured to convey address packets specifying a particular coherency unit in broadcast mode. The address network included in a different one of the nodes may be configured to convey address packets specifying that coherency unit in point-to-point mode.

PRIORITY INFORMATION

[0001] This application claims priority to U.S. provisional application serial No. 60/461,997, entitled "MULTI-NODE COMPUTER SYSTEM In WHICH NETWORKS IN DIFFERENT NODES IMPLEMENT DIFFERENT CONVEYANCE MODES", filed Apr. 11, 2003.

h e b b g eeef c e b e ge

Previous Doc Next Doc Go to Doc#

h e b b g e e e f c e b e ge

L7: Entry 11 of 49 File: PGPB Jan 17, 2002

PGPUB-DOCUMENT-NUMBER: 20020007443
PGPUB-FILING-TYPE: new
DOCUMENT-IDENTIFIER: US 20020007443 A1

TITLE: Scalable multiprocessor system and cache coherence method

PUBLICATION-DATE: January 17, 2002

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-47
Gharachorloo, Kourosh	Menlo Park	CA	US	
Barroso, Luiz A.	Mountain View	CA	US	
Ravishankar, Mosur K.	Mountain View	CA	US	
Stets, Robert J. JR.	Palo Alto	CA	US	
Scales, Daniel J.	Mountain View	CA	US	

APPL-NO: 09/ 878982 [PALM]
DATE FILED: June 11, 2001

RELATED-US-APPL-DATA:
Application is a non-provisional-of-provisional application 60/210675, filed June 10, 2000,

FOREIGN-APPL-PRIORITY-DATA:

COUNTRY	APPL-NO	DOC-ID	APPL-DATE
US	60210675	2000US-60210675	June 10, 2000

INT-CL: [07] G06 F 12/00

US-CL-PUBLISHED: 711/141; 711/117
US-CL-CURRENT: 711/141; 711/117

REPRESENTATIVE-FIGURES: 3

ABSTRACT:

The present invention relates generally to multiprocessor computer system, and particularly to a multiprocessor system designed to be highly scalable, using efficient cache coherence logic and methodologies. More specifically, the present invention is a system and method including a plurality of processor nodes configured to execute a cache coherence protocol that avoids the use of negative acknowledgment messages (NAKs) and ordering requirements on the underlying transaction-message interconnect/network and services most 3-hop transactions with only a single visit to the home node.

h e b b g eeef c e b e ge

RELATED APPLICATIONS

[0001] This application is related to the following U.S. patent applications:

[0002] System and Method for Daisy Chaining Cache Invalidation Requests in a Shared-memory Multiprocessor System, filed Jun. 11, 2001, attorney docket number 9772-0329-999; and

[0003] Multiprocessor Cache Coherence System and Method in Which Processor Nodes and Input/Output Nodes Are Equal Participants, filed Jun. 11, 2001, attorney docket number 9772-0324-999; and

[0004] Cache Coherence Protocol Engine And Method For Processing Memory Transaction in Distinct Address Subsets During Interleaved Time Periods in a Multiprocessor System, filed Jun. 11, 2001, attorney docket number 9772-0327-999.

Previous Doc Next Doc Go to Doc#

L7: Entry 21 of 49 File: USPT Jan 6, 2004

US-PAT-NO: 6675265
DOCUMENT-IDENTIFIER: US 6675265 B2

TITLE: Multiprocessor cache coherence system and method in which processor nodes and input/output nodes are equal participants

DATE-ISSUED: January 6, 2004

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Barroso; Luiz A.	Mountain View	CA		
Gharachorloo; Kourosh	Menlo Park	CA		
Nowatzyk; Andreas	San Jose	CA		
Ravishankar; Mosur K.	Mountain View	CA		
Stets, Jr.; Robert J.	Palo Alto	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Hewlett-Packard Development Company, L.P.	Houston	TX			02

APPL-NO: 09/ 878984 [PALM]
DATE FILED: June 11, 2001

PARENT-CASE:
RELATED APPLICATIONS This application claims the benefit of Provisional Application No. 60/210,675, filed Jun. 10, 2000. This application is related to, and hereby incorporates by reference, the following U.S. patent applications: Scalable Multiprocessor System And Cache Coherence Method, filed Jun. 11, 2001, Ser. No. 09/878,982. System And method for Daisy Chaining Cache Invalidation Requests In A Shared-Memory Multiprocessor System, filed Jun. 11, 2001, Ser. No. 09/878,955. Cache Coherence Protocol Engine And Method For Processing Memory Transaction ln Distinct Address Subsets During Interleaved Time Periods in A Multiprocessor System, tiled Jun. 11, 2001, Ser. No. 09/878,983. The present invention relates generally to multiprocessor computer system, and particularly to a multiprocessor system designed to be highly scalable, using efficient cache coherence logic and methodologies.

INT-CL: [07] G06 F 12/00

US-CL-ISSUED: 711/141; 711/144, 711/145
US-CL-CURRENT: 711/141; 711/144, 711/145

FIELD-OF-SEARCH: 711/141, 711/117, 711/118, 711/142, 711/143, 711/144, 711/145, 711/121, 711/147

PRIOR-ART-DISCLOSED:

h e b b g eeef c e b e ge

U.S. PATENT DOCUMENTS

Search Selected | Search ALL | Clear

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	5634110	May 1997	Laudon et al.	711/145
☐	5963975	October 1999	Boyle et al.	711/147
☐	6263403	July 2001	Traynor	711/133
☐	6438653	August 2002	Akashi et al.	711/128
☐	6493809	December 2002	Safranek et al.	711/167

ART-UNIT: 2187

PRIMARY-EXAMINER: Elmore; Reba I.

ASSISTANT-EXAMINER: Takeguchi; Kathy

ABSTRACT:

A computer system has a plurality of processor nodes and a plurality of input/output nodes. Each processor node includes a multiplicity of processor cores, an interface to a local memory system and a protocol engine implementing a predefined cache coherence protocol. Each processor core has an associated memory cache for caching memory lines of information. Each input/output node includes no processor cores, an input/output interface for interfacing to an input/output bus or input/output device, a memory cache for caching memory lines of information and an interface to a local memory subsystem. The local memory subsystem of each processor node and input/output node stores a multiplicity of memory lines of information. The protocol engine of each processor node and input/output node implements the same predefined cache coherence protocol.

14 Claims, 30 Drawing figures

Previous Doc Next Doc Go to Doc#

h e b b g eeef c e b e ge

L7: Entry 38 of 49 File: USPT Jun 12, 2001

DOCUMENT-IDENTIFIER: US 6247091 B1
** See image for Certificate of Correction **
TITLE: Method and system for communicating interrupts between nodes of a multinode computer system

Brief Summary Text (5):
These computers may be classified by how they share information among the processors. Shared-memory multiprocessor computers offer a common memory address space that all processors can access. Processes within a program communicate through shared variables in memory which allow them to read or write to the same memory location in the computer. Message passing multiprocessor computers, on the other hand, have a separate memory space for each processor. Processes communicate through messages to each other.

Brief Summary Text (7):
Multiprocessor computers with distributed shared memory are often organized into nodes with one or more processors per node. Also included in the node are local memory for the processors, a remote cache for caching data obtained from memory in other nodes, and logic for linking the node with other nodes in the computer. A processor in a node communicates directly with the local memory and communicates indirectly with memory on other nodes through the remote cache. For example, if the desired data is in local memory, a processor obtains the data directly from local memory. But if the desired data is stored in memory in another node, the processor must access its remote cache to obtain the data. A cache hit occurs if the data has been obtained recently and is presently stored in the cache. Otherwise a cache miss occurs, and the cache must obtain the desired data from the local memory in another node through the linking logic.

Brief Summary Text (10):
Bus-based interrupt schemes, however, cannot communicate interrupts across the network of a multinode multiprocessor system because the nodes are not connected by a bus. (The difference between a bus and a network is well defined. See, for example, "Interconnection Networks," Computer Architecture A Quantitative Approach, .sub.2 nd Ed. (1996).) Instead, a second interrupt mechanism must be added to handle interrupts sent via the network from a processor on one node to a processor on another node. The obvious solution is to treat an interrupt like data and provide an interrupt register with a memory address in each node. A requesting processor in one node then interrupts a processor in a second node by writing an interrupt request to the address of the interrupt register in the second node. The request is then sent by way of the network to the second node. Hardware in the second node reads the interrupt register and places the interrupt request on the second node's bus for the second processor to read.

Detailed Description Text (3):
FIG. 1 is a block diagram of a multinode, multiprocessor computer system 10 in accordance with the invention. System 10 uses a computer architecture based on Distributed-Shared Memory (DSM). Four nodes 12-18 are shown connected by a system interconnect 20 that permits any node to communicate with any other node. Specifically, the purpose of interconnect 20 is to allow processors in any node to access the memory resident in any other node with cache coherency guaranteed.

h e b b g e e e f c e b e g

System interconnect 20 is a switch-based interconnection network that uses the Scalable Coherent Interface (SCI) interconnection mechanism. SCI is an IEEE-approved standard, or protocol (1596), well documented in a number of publications including IEEE Std 1596-1992 (Aug. 2, 1993) and Multiprocessor interconnection using SCI, a Master Thesis by Ivan Tving, DTH ID-E 579 (1994), both of which are hereby incorporated by reference.

Current US Cross Reference Classification (4):
709/251

Previous Doc Next Doc Go to Doc#

h e b b g e e e f c e b e g

L7: Entry 38 of 49 File: USPT Jun 12, 2001

US-PAT-NO: 6247091
DOCUMENT-IDENTIFIER: US 6247091 B1
** See image for Certificate of Correction **

TITLE: Method and system for communicating interrupts between nodes of a multinode computer system

DATE-ISSUED: June 12, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Lovett; Thomas D.	Portland	OR		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
International Business Machines Corporation	Armonk	NY			02

APPL-NO: 08/ 848545 [PALM]
DATE FILED: April 28, 1997

INT-CL: [07] G06 F 13/24

US-CL-ISSUED: 710/260; 710/266, 710/261, 710/263, 710/264, 710/267, 710/268, 709/251, 709/249, 709/253, 709/230, 370/402
US-CL-CURRENT: 710/260; 370/402, 709/230, 709/249, 709/251, 709/253, 710/261, 710/263, 710/264, 710/266, 710/267, 710/268

FIELD-OF-SEARCH: 710/260, 710/262-264, 710/266, 710/268, 710/269, 709/745, 709/238, 709/212, 709/249, 709/251, 709/250, 709/230, 370/392, 370/402

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	4604500	August 1986	Brown et al.	379/269
☐	4768149	August 1988	Konopik et al.	395/867
☐	5109522	April 1992	Lent et al.	395/500
☐	5283904	February 1994	Carson et al.	395/739
☐	5369748	November 1994	McFarland et al.	711/118

h e b b g eeef c e b e ge

☐	5428799	June 1995	Woods et al.	710/266
☐	5481725	January 1996	Jayakumar et al.	710/48
☐	5511200	April 1996	Jayakumar	710/266
☐	5566171	October 1996	Levinson	370/352
☐	5598541	January 1997	Malladi	710/106
☐	5623494	April 1997	Rostoker et al.	370/397
☐	5625563	April 1997	Rostoker et al.	364/488
☐	5638518	June 1997	Malladi	709/251
☐	5640399	June 1997	Rostoker et al.	370/392
☐	5678057	October 1997	Rostoker et al.	712/11
☐	5706514	January 1998	Bonola	709/104
☐	5715274	February 1998	Rostoker et al.	375/200
☐	5742843	April 1998	Koyanagi et al.	712/14
☐	5790530	August 1998	Moh et al.	370/363
☐	5832279	November 1998	Rostoker et al.	710/266
☐	5892956	April 1999	Qureshi et al.	710/260
☐	5944798	August 1999	McCarty et al.	709/251

OTHER PUBLICATIONS

"Design of the APIC: A High Performance ATM Host Hetwork Interface Chip", Dittia et al., 1995 IEEE.*
"The APIC Approach to High Performance Network Interface Design: Protected and Other Techniques", Dittia et al., 1997 IEEE.*
"Efficient Interprocessor Communications in a Tightly-Coupled Homogeneous Multiprocessor System", van der Wal et al. 1990 IEEE.*
MultiProcessor Specification, Intel, ver. 1.4, Jul. 1, 1995, Rev. Aug. 1996 (appendix E added)

ART-UNIT: 213

PRIMARY-EXAMINER: Teska; Kevin J.

ASSISTANT-EXAMINER: Thomson; William

ATTY-AGENT-FIRM: Klarquist Sparkman Campbell Leigh and Whinston LLP

ABSTRACT:

Each node of multinode computer system includes an interrupt controller, a pair of send and receive queues, and a state machine for communicating interrupts between nodes. The communication among the interrupt controller, the state machine, and the queues is coordinated by a queue manager. For sending an interrupt, the interrupt controller accepts an interrupt placed on a bus within the node and intended for another node and stores it in the send queue. The controller then notifies the interrupt source that the interrupt has been accepted before it is transmitted to other node. The interrupt has a first form suitable for transmission on the bus. A state machine within the node takes the interrupt from the send queue and puts the interrupt into a second form suitable for transmission across a network connecting

the multiple nodes. For receiving an interrupt, the state machine accepts an interrupt from another node and stores it in the receive queue, notifying the interrupt source that the interrupt has been accepted before its is placed on the node bus. The interrupt has the second form suitable for transmission across the network. The interrupt controller takes the interrupt from the receive queue and puts it in the first form suitable for transmission on the bus.

20 Claims, 12 Drawing figures

Previous Doc Next Doc Go to Doc#

h e b b g e e e f c e b e ge

L7: Entry 47 of 49 File: USPT Jul 14, 1998

US-PAT-NO: 5781757
DOCUMENT-IDENTIFIER: US 5781757 A

TITLE: Adaptive scalable cache coherence network for a multiprocessor data processing system

DATE-ISSUED: July 14, 1998

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Deshpande; Sanjay Raghunath	Austin	TX		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
International Business Machines Corporation	Armonk	NY			02

APPL-NO: 08/ 747587 [PALM]
DATE FILED: November 13, 1996

PARENT-CASE:
This is a continuation of application Ser. No. 08/320,484, filed 11 Oct. 1994 now abandoned.

INT-CL: [06] G06 F 13/00

US-CL-ISSUED: 395/473; 395/200.02, 395/200.03, 395/200.1, 395/200.16, 395/200.15, 395/200.21, 395/297, 395/300, 395/446 , 395/447, 395/448, 395/468, 395/449, 395/471, 395/472, 395/473, 395/730, 395/731, 395/800
US-CL-CURRENT: 711/146; 709/201, 710/117, 710/120, 710/242, 710/243, 711/119, 711/120, 711/121, 711/122, 711/141, 711/144, 711/145, 712/28, 712/30

FIELD-OF-SEARCH: 395/447, 395/446, 395/448, 395/468, 395/449, 395/473, 395/472, 395/471, 395/200.16, 395/297, 395/200.21, 395/200.15, 395/300, 395/730, 395/731, 395/200.02

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	4049974	September 1977	Boone et al.	326/95

h e b b g eeef c e b e ge

☐	4367420	January 1983	Foss et al.	326/119
☐	4570084	February 1986	Griffin et al.	326/98
☐	4755930	July 1988	Wilson, Jr. et al.	395/449
☐	4794342	December 1988	Kimura	330/2
☐	4896057	January 1990	Yang et al.	326/116
☐	4907184	March 1990	Nakano et al.	369/736
☐	5025365	June 1991	Mathur et al.	395/448
☐	5030857	July 1991	Sanwo et al.	326/86
☐	5250860	October 1993	Chu et al.	326/60
☐	5265235	November 1993	Sindhu et al.	395/447
☐	5274782	December 1993	Chalasani et al.	395/325
☐	5313609	May 1994	Baylor et al.	395/448
☐	5345578	September 1994	Manasse	395/473
☐	5440698	August 1995	Sindhu et al.	395/200.08
☐	5530646	June 1996	Schoen	364/421

OTHER PUBLICATIONS

Lipovski, G. J., An Organization for Optical Linkages Between Integrated Circuits, National Computer Conference, 1977, pp. 227-236.
Dae-Wha Seo, et al, Directory-Based Cache Coherence Scheme Using Number-Balanced Binary Tree, Microprocess. Microprogr. (Netherlands), vol. 37, No. 1-5, pp. 37-40, Jan. 1993, 4 REF.
Farrens, M., et al., CCHIME: A Cache Coherent Hybrid Interconnected Memory Extension, Proceedings: Sixth International Parallel Processing Symposium, (Cat. No. 92TH0419-2), IEEE Computer Soc. Press, xviii+693 pp., pp. 573-577, 1992, 11 REF.
Gupta, S., et al., Stanford DASH Multiprocessor: The Hardware and Software Approach, Proceedings: 4th International Parallel Architectures and Languages Europe Conference, 1992, Springer-Verlag, xvii+894 pp., pp. 802-805, 1992 40 REF.
Michael, W., Directory-Based Coherency Protocol For a Ring-Connected Multiprocessor-Array, Comput. Archit. News (USA), vol. 20, No. 2, p. 437, May 1992, 0 REF.
Al-Sadoun, H.B., et al., Cache Coherency In Multiple Bus Systems, Int. J. Electron (UK), vol. 73, No. 3, pp. 497-522, Sept. 1992, 33 REF.
Tamir, Y., et al., Hierarchical Coherency Management For Shared Virtual Memory Multicomputers, J. Parallel Distribut. Comput. (USA), vol. 15, No. 4, pp. 408-419, Aug. 1992, 17 REF.
Thapar, M., et al., Scalable Cache Coherence For Shared Memory Multiprocessors, Proceedings: Parallel Computation, First International ACPC Conference, Springer-Verlag, ix+451 pp., pp. 1-12, 1992, 13 REF.
Dahlgren, F., et al., Reducing Write Latencies For Shared Data In A Multiprocessor With A Multistage Network, Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences (Cat. No. 91TH0394-7), IEEE Comput. Soc. Press, 4 vol., (xv+831+877+xii+670+xiii+729) pp., pp. 449-456, vol. 1, 1991, 14 REF.
Tiruveedhula, V., et al., Performance Analysis Of A Cache-Based Multiprocessor System Using Data Movement Policies, Microprocess. Microprogr., (Netherlands), vol. 33, No. 4, pp. 237-248, Jun. 1992, 7 REF.
Marquardt, D.E., et al., A Cache-Coherent, Distributed Memory Multiprocessor System and Its Performance Analysis, IEICE Trans. Inf. Syst. (Japan), vol. E75-D, No. 3,

pp. 274-290, May 1992, 34 REF.
Yang, Q., et al., Design Of An Adaptive Cache Coherence Protocol For Large Scale Multiprocessors, IEEE Trans. Parallel Distrib Syst. (USA), vol. 3, No. 3, pp. 281-293; May 1992, 27 REF.
Yang, Q., et al., An Adaptive Cache Coherence Scheme For Hierarchical Shared-Memory Multiprocessors, Proceedings of the Second IEEE Symposium on Parallel and Distributed Processing 1990 (Cat. No. TH0328-5), IEEE Comput. Soc. Press, xix+892 pp. 318-325, 1990, 22 REF.
Mori, S., et al., The Kyushu University Reconfigurable Parallel Processor-Cache Architecture and Cache Coherence Schemes, Proceedings of the International Symposium on Shared Memory Multiprocessing, Inf. Process. Soc. Japan, 251 pp., pp. 218-219, 1991, 23 REF.
Thapar, M. et al., Cache Coherence For Large Scale Shared Memory Multiprocessors, Comput. Archit. News (USA), vol. 19, No. 1, pp. 114-119, March 1991, 12 REF.
Thaper, M., et al., Scalable Cache Coherence for Large Shared Memory Multiprocessors, Proceedings: 1990 Joint International Conference on Vector and Parallel Processing, Springer-Verlag, xi+900 pp., pp. 592-603, 1990, 19 REF.
Lenoski, D., et al., Design of Scalable Shared-Memory Multiprocessors: The DASH Approach, COMPCON Spring '90, Thirty-Fifth IEEE Computer Society International Conference, Intellectual Leverage, Digest of Papers, (Cat. No. 90CH2843-1), IEEE Comput Soc., xvi+644 pp., pp. 62-67, 1990, 11 REF.
Yang, Q., et al., Performance Analysis Of A Cache-Coherent Multiprocessor Based On Hierarchical Multiple Buses, Proceedings of the International Conference on Databases, Parallel Architectures and Their Applications, 1990, (Cat. No. 90CH2728-4), IEEE Comput. Soc. Press, xiv+570 pp., pp. 248-257, 1990, 24 REF.
Marquardt, D.E., et al., C/Sup 2/MP: A Cache-Coherent, Distributed Memory Multiprocessor-System, Proceedings of Supercomputing '89, ACM, xviii+894 pp., pp. 466-475, 1989, 39 REF.
Bhuyan, L.N. et al., Analysis of MIN Based Multiprocessors With Private Cache Memories, Proceedings of the 1989 International Conference on Parallel Processing, (Cat. No. 89CH2701-1), Pennsylvania State Univ. Press, 3 vol., (xv+137+xiii+263+xiii+262) pp., pp. 51-58 vol. 1, 1989, 20 REF.
Owicki, S., et al., Evaluating The Performance Of Software Cache Coherence, ASPLOS-III Proceedings. Third International Conference on Architectural Support For Programming Languages and Operating Systems, ACM x+303pp., pp. 230-242, 1989, 33 REF.
Winsor, D.C., et al., Analysis Of Bus Hierarchies For Multiprocessors15th Annual International Symposium on Computer Architecture, Conference Proceedings (Cat. No. 88CH2545-2), IEEE Comput. Soc. Press, xi+461 pp., pp. 100-107, 1988, 20 REF.
Bhuyan, L.N., et al., Multistage Bus Network (MBN) An Interconnection Network For Cache Coherent Multiprocessors, Proceedings of the Third IEEE Symposium on Parallel and Distributed Processing (Cat. No. 91TH0396-2), IEEE Comput. Soc. Press, xvi+903 pp., pp. 780-787, 1991, 14 REF.
Ganapathy, K.N., et al., Yield Optimization In Large RAM's With Hierarchical Redundancy, Special Brief Papers, IEEE Journal of Solid-State Circuits, vol. 26, No. 9, Sep. 1991.
Deshpande, Sanjay, et al., Scalability Of A Binary Tree On A Hypercube, Department of Electrical and Computer Engineering, University of Texas at Austin, 1986.

ART-UNIT: 238

PRIMARY-EXAMINER: Swann; Tod R.

ASSISTANT-EXAMINER: Tran; Denise

ATTY-AGENT-FIRM: Henkler; Richard A. Russell; Brian F. Dillon; Andrew J.

ABSTRACT:

h e b b g eeef c e b e ge

A cache coherence network for transferring coherence messages between processor caches in a multiprocessor data processing system is provided. The network includes a plurality of processor caches associated with a plurality of processors, and a binary logic tree circuit which can separately adapt each branch of the tree from a broadcast configuration during low levels of coherence traffic to a ring configuration during high levels of coherence traffic. A cache snoop-in input receives coherence messages and a snoop-out output outputs, at the most, one coherence message per current cycle of the network timing. A forward signal on a forward output indicates that the associated cache is outputting a message on snoop-out during the current cycle. A cache outputs received messages in a queue on the snoop-out output, after determining any response message based on the received message. The binary logic tree circuit has a plurality of binary nodes connected in a binary tree structure. Each branch node has a snoop-in, a snoop-out, and a forward connected to each of a next higher level node and two lower level nodes. A forward signal on a forward output indicates that the associated node is outputting a message on snoop-out to the higher node during the current cycle. Each branch ends with multiple connections to a cache at the cache's snoop-in input, snoop-out output, and forward output.

7 Claims, 11 Drawing figures

Previous Doc Next Doc Go to Doc#

h e b b g e e e f c e b e ge

First Hit Fwd Refs Previous Doc Next Doc Go to Doc#

End of Result Set

Generate Collection Print

L8: Entry 1 of 1 File: USPT Oct 21, 2003

US-PAT-NO: 6636926
DOCUMENT-IDENTIFIER: US 6636926 B2

TITLE: Shared memory multiprocessor performing cache coherence control and node controller therefor

DATE-ISSUED: October 21, 2003

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Yasuda; Yoshiko	Tokorozawa			JP
Hamanaka; Naoki	Tokyo			JP
Shonai; Toru	Hachioji			JP
Akashi; Hideya	Kunitachi			JP
Tsushima; Yuji	Kokubunji			JP
Uehara; Keitaro	Kokubunji			JP

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Hitachi, Ltd.	Tokyo			JP	03

APPL-NO: 09/ 740816 [PALM]
DATE FILED: December 21, 2000

FOREIGN-APPL-PRIORITY-DATA:

COUNTRY	APPL-NO	APPL-DATE
JP	11-366235	December 24, 1999

INT-CL: [07] G06 F 13/00, G06 F 15/167

US-CL-ISSUED: 710/305; 710/317, 711/141, 709/213, 700/5
US-CL-CURRENT: 710/305; 700/5, 709/213, 710/317, 711/141

FIELD-OF-SEARCH: 710/305, 710/317, 710/300, 710/62, 710/4, 710/72, 711/141, 711/148, 711/120, 709/213, 709/214, 709/251, 700/5, 712/14, 712/211

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

h e b b g eeef c e b e ge

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	4747043	May 1988	Rodman	
☐	6011791	January 2000	Okada et al.	
☐	6092173	July 2000	Sasaki et al.	
☐	6378029	April 2002	Venkitakrishnan et al.	
☐	6466825	October 2002	Wang et al.	

OTHER PUBLICATIONS

"RISC System/6000SMP System," 1995 Comcon95 Proceedings, pp. 102-109.
"Starfire: Extending the SMP Envelope," 1998 Micro Jan./Feb. pp. 39-49.

ART-UNIT: 2181

PRIMARY-EXAMINER: Ray; Gopal C.

ABSTRACT:

Each node includes a node controller for decoding the control information and the address information for the access request issued by a processor or an I/O device, generating, based on the result of decoding, the cache coherence control information indicating whether the cache coherence control is required or not, the node information and the unit information for the transfer destination, and adding these information to the access request. An intra-node connection circuit for connecting the units in the node controller holds the cache coherence control information, the node information and the unit information added to the access request. When the cache coherence control information indicates that the cache coherence control is not required and the node information indicates the local node, then the intra-node connection circuit transfers the access request not to the inter-node connection circuit interconnecting the nodes but directly to the unit designated by the unit information.

19 Claims, 16 Drawing figures

Previous Doc Next Doc Go to Doc#

US006343346B1

(12) **United States Patent**
Olnowich

(10) **Patent No.: US 6,343,346 B1**
(45) **Date of Patent: Jan. 29, 2002**

(54) **CACHE COHERENT NETWORK ADAPTER FOR SCALABLE SHARED MEMORY PROCESSING SYSTEMS**

(75) Inventor: **Howard Thomas Olnowich**, Endwell, NY (US)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: 09/516,393

(22) Filed: Mar. 1, 2000

Related U.S. Application Data

(62) Division of application No. 08/891,404, filed on Jul. 10, 1997, now Pat. No. 6,092,155.

(51) Int. Cl.[7] G06F 12/08

(52) U.S. Cl. 711/142; 711/121; 711/122; 711/153

(58) Field of Search 711/121, 122, 711/143, 144, 145, 146, 148, 153, 142; 709/214, 215, 218, 250

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,399,504	A	8/1983	Watts et al.	364/200
4,562,539	A	12/1985	Vince	395/200.08
4,755,930	A	7/1988	Wilson, Jr. et al.	395/468
4,811,216	A	3/1989	Bishop et al.	364/228.3
4,965,719	A	10/1990	Shoens et al.	364/200
5,123,106	A	6/1992	Otsuki et al.	711/153
5,218,688	A	6/1993	Nishida	711/151
5,303,362	A	4/1994	Butts, Jr. et al.	711/121
5,313,609	A	5/1994	Baylor et al.	395/468
5,394,555	A	2/1995	Hunter et al.	711/148
5,442,758	A	8/1995	Slingwine et al.	395/508
5,444,705	A	8/1995	Olnowich et al.	370/85.6
5,452,447	A	9/1995	Nelson et al.	395/200.08
5,499,349	A	3/1996	Nikhil et al.	395/375
5,522,058	A	5/1996	Iwasa et al.	711/145
5,530,816	A	6/1996	Holt	395/375
5,535,116	A	7/1996	Gupta et al.	395/200.08
5,537,569	A	7/1996	Masubuchi	395/200.08
5,537,574	A	7/1996	Elko et al.	395/200.08
5,557,792	A	9/1996	Jones et al.	395/200.08
5,561,809	A	10/1996	Elko et al.	395/200.08
5,592,625	A	1/1997	Sandberg	395/200.08
5,610,953	A	3/1997	Olnowich et al.	375/373
5,611,049	A	3/1997	Pitts	395/200.08
5,864,671	A *	1/1999	Hagersten et al.	709/213
5,987,506	A *	11/1999	Carter et al.	709/213

(List continued on next page.)

OTHER PUBLICATIONS

M. Dubois et al. "Effects of Cache Coherency in Multiprocessors", *IEEE Transactions on Computers*, vol. C-31, No. 11, Nov. 1982.

Primary Examiner—Hiep T. Nguyen
(74) *Attorney, Agent, or Firm*—Shelley M Beckstrand

(57) **ABSTRACT**

A shared memory parallel processing system interconnected by a multi-stage network combines new system configuration techniques with special-purpose hardware to provide remote memory accesses across the network, while controlling cache coherency efficiently across the network. The system configuration techniques include a systematic method for partitioning and controlling the memory in relation to local verses remote accesses and changeable verses unchangeable data. Most of the special-purpose hardware is implemented in the memory controller and network adapter, which implements three send FIFOs and three receive FIFOs at each node to segregate and handle efficiently invalidate functions, remote stores, and remote accesses requiring cache coherency. The segregation of these three functions into different send and receive FIFOs greatly facilitates the cache coherency function over the network. In addition, the network itself is tailored to provide the best efficiency for remote accesses.

61 Claims, 41 Drawing Sheets

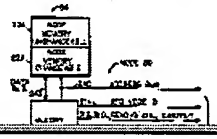

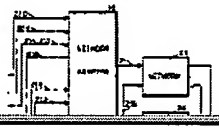

US006044438A

United States Patent [19]

Olnowich

[11] Patent Number: 6,044,438

[45] Date of Patent: Mar. 28, 2000

[54] **MEMORY CONTROLLER FOR CONTROLLING MEMORY ACCESSES ACROSS NETWORKS IN DISTRIBUTED SHARED MEMORY PROCESSING SYSTEMS**

[75] Inventor: **Howard Thomas Olnowich**, Endwell, N.Y.

[73] Assignee: **International Business Machines Corporation**, Armonk, N.Y.

[21] Appl. No.: 08/890,341

[22] Filed: Jul. 10, 1997

[51] Int. Cl.[7] G06F 13/14

[52] U.S. Cl. 711/130; 711/120; 711/141; 711/150; 707/201; 709/213; 709/214

[58] Field of Search 711/120, 130, 711/141, 150; 707/201; 709/213, 214, 250

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,399,504	8/1983	Watts et al. .	
4,562,539	12/1985	Vince .	
4,755,930	7/1988	Wilson, Jr. et al. .	
4,965,719	10/1990	Shoens et al. .	
5,313,609	5/1994	Baylor et al. .	
5,442,758	8/1995	Slingwine et al. .	
5,444,705	8/1995	Olnowich et al. .	
5,452,447	9/1995	Nelson et al. .	
5,499,349	3/1996	Nikhil et al. .	
5,530,816	6/1996	Holt .	
5,535,116	7/1996	Gupta et al. .	
5,537,569	7/1996	Masubuchi .	
5,537,574	7/1996	Elko et al. .	
5,557,792	9/1996	Josten et al. .	
5,561,809	10/1996	Elko et al. .	
5,592,625	1/1997	Sandberg .	
5,610,953	3/1997	Olnowich et al. .	
5,611,049	3/1997	Pitts .	
5,737,568	4/1998	Hamaguchi et al. .	
5,832,534	11/1998	Singh et al.	711/141

OTHER PUBLICATIONS

M. Dubcise et al. "Effects of Cache Coherency in Multiprocessors", *IEEE Transactions on Computers*, vol. C-31, No. 11, Nov. 1982.

Primary Examiner—Kriana Lim
Attorney, Agent, or Firm—Shelley M Beckstrand

[57] **ABSTRACT**

A shared memory parallel processing system interconnected by a multi-stage network combines new system configuration techniques with special-purpose hardware to provide remote memory accesses across the network, while controlling cache coherency efficiently across the network. The system configuration techniques include a systematic method for partitioning and controlling the memory in relation to local verses remote accesses and changeable verses unchangeable data. Most of the special-purpose hardware is implemented in the memory controller and network adapter, which implements three send FIFOs and three receive FIFOs at each node to segregate and handle efficiently invalidate functions, remote stores, and remote accesses requiring cache coherency. The segregation of these three functions into different send and receive FIFOs greatly facilitates the cache coherency function over the network. In addition, the network itself is tailored to provide the best efficiency for remote accesses.

16 Claims, 41 Drawing Sheets